

ACTION RECOGNITION IN STILL IMAGES BY LATENT SUPERPIXEL CLASSIFICATION

Shaukat Abidi, Massimo Piccardi, Mary-Anne Williams

Faculty of Engineering and IT, University of Technology, Sydney

ABSTRACT

Action recognition from still images is an important task of computer vision applications such as image annotation, robotic navigation, video surveillance and several others. Existing approaches mainly rely on either bag-of-feature representations or articulated body-part models. However, the relationship between the action and the image segments is still substantially unexplored. For this reason, in this paper we propose to approach action recognition by leveraging an intermediate layer of “superpixels” whose latent classes can act as attributes of the action. In the proposed approach, the action class is predicted by a structural model (learnt by Latent Structural SVM) based on measurements from the image superpixels and their latent classes. Experimental results over the challenging Stanford 40 Actions dataset report a significant average accuracy of 74.06% for the positive class and 88.50% for the negative class, giving evidence to the performance of the proposed approach¹.

Index Terms— Action recognition from still images, superpixel segmentation, latent structural SVM.

1. INTRODUCTION AND RELATED WORK

Automated recognition of actions in still images can play an important role for annotation of image catalogues, including the large collections of images which are increasingly made available by social networks. Actions which can be plausibly recognised from still imagery are those inferrable from the actors’ poses and the presence of relevant objects: examples range from “taking a picture” and “having a barbecue”, to “throwing a javelin” or “playing guitar”. Moreover, recognition from single frames could also prove of fundamental value for recognising actions in video. For instance, in surveillance videos it is not uncommon to clearly sight an actor for only a few frames due to repeated occlusions. In such cases, it is not easy to recognise the action in dynamical terms, i.e., as the temporal evolution of a measurement vector. Rather, inference must be obtained as the cumulative evidence from a (possibly small) set of individual frames. Also in robotics, the varying camera viewpoint may make it easier to recognise actions from isolated frames than from sequences. Estimation




Fig. 1. Action recognition: bottom layer: superpixel segmentation and feature extraction; intermediate layer: superpixel classification; top variable: action class.

from still frames is therefore a foundational technology for all these cases.

The most straightforward solution to recognise actions from still images is to compute a bag-of-feature representation of the image and use it for classifying it into a relevant set of action classes [1, 2, 3]. Useful features include local texture descriptors such as the histogram of oriented gradients (HOG), dense SIFT, GIST and several others [4, 5, 6]. Bag-of-features analysis usually discards the spatial coordinates at which descriptors are collected since it focusses on textural rather than spatial or structural information. These approaches have reported very interesting results on challenging still image action datasets such as those described in [7, 8, 9]. At the opposite end of the spectrum are approaches based on the explicit recovery of body parts and the incorporation of structural information in the recognition process [10, 11]. The baseline model is a latent part-based model akin to Pictorial Structure which can be estimated as a joint, conditional or max-margin model [12, 13]. Delaitre *et al.* has reported a comparison between a bag-of-features and a structural approach, showing that hybridisation of the

¹ ©2015 IEEE

two can be a way to capture the benefits of both models [1]. A recent survey from Guo and Lai offers a comprehensive outline of the research in this area [14].

Overall, it appears that the existing approaches have not substantially explored the underlying relationship between the action class and the segments of the containing image. For this reason, in this paper we propose to approach action recognition in still images by leveraging the latent classes of the image’s “superpixels” (homogenous regions obtained from over-segmentation of the image [15]). To this aim, we have designed a graphical model with the action as its root node and a fully-connected layer of superpixel classes to capture the relationship with the image segments (Fig. 1). A rich measurement vector is extracted from each superpixel, and model training is provided by a latent structural SVM approach.

The rest of the paper is organised as follows: Section 2 describes the graphical model for action recognition and the detectors for the superpixels’ classes. Section 3 overviews the latent structural SVM framework. Section 4 describes the experiments and discusses results. Section 5 highlights conclusions and future work.

2. ACTION RECOGNITION BY SUPERPIXEL CLASSIFICATION

The fundamental step taken in our approach is the decomposition of an image into small and coherent patches commonly referred to as superpixels. Our underlying assumption is that certain actions can be recognised effectively from an image by utilising useful information from the superpixels. Given their homogeneous nature, superpixels can be assigned single class labels of the type of “sky”, “road”, “face” and others, leading to a form of image (over-)segmentation (Fig. 1, bottom layers). While each superpixel in an image can be classified individually using a trained classifier, a recent paper from Pei *et al.* [16] has shown that superpixel classification proves more accurate if all the superpixels in the image are classified jointly. Accordingly, our approach consists of two main stages: in the first stage, we pre-train a superpixel classifier, or a set of object detectors, from a supervised set of image regions and we use it to compute class scores for all superpixels in a given image. In the second stage, such scores are used as measurements in a graphical model that provides optimal, joint decisions for all superpixels and the action class. In the rest of this section, we describe the graphical model and the object detectors.

2.1. The graphical model

The proposed graphical model comprises three sets of variables, namely measurements (x), hidden nodes, or states, (h) and an output node (y). The measurements are a vector of detector scores for each superpixel, the hidden nodes are their

classes, and the output node is the categorical variable for the action. The nodes are connected by three different types of edges: a) edges connecting measurements and states, b) edges over state pairs, and c) edges between states and the action class. Given that the dependencies between states may extend over the entire image, we assume that the states and their edges form a fully-connected graph. Approximate inference is provided by a greedy algorithm that iteratively maximises over each state in turn.

Noting the number of superpixels in the i -th image in a training set as T_i , we have $x_i = [x_i^1, \dots, x_i^t, \dots, x_i^{T_i}]$, with x_i^t a D -dimensional vector of scores; $h_i = [h_i^1, \dots, h_i^t, \dots, h_i^{T_i}]$, with $h_i^t \in \{1, \dots, K\}$ a superpixel class; and $y_i \in \{0, 1\}$ a binary variable for a given action class. We build one such model for each action class. At training time, the action class is supervised while states are hidden.

2.2. Object detectors

The first step of our processing pipeline is the decomposition of the image into superpixels. For this task, we use the efficient graph-based segmentation algorithm proposed by Felzenszwalb *et al.* [15] that was also adopted by [16]. This step achieves good over-segmentation of the image into regions of predominantly homogeneous nature, and errors in this process can be tolerated by the ensuing soft-assignment stage. Fig. 1 shows an example of superpixel segmentation.

To classify the superpixels, we have used the class set of the MSRC-21 dataset consisting of 23 diverse classes from typical background and foreground objects [17]. Similarly to [16], we use a combination of appearance-based and bag-of-features descriptors as the feature vector. The appearance-based descriptor has 51 features, comprising of: 1) 40 color features measuring mean, standard deviation, skewness and kurtosis of RGB, LAB, YCrCb color space channels and the gray image; and 2) 11 texture features obtained from the application of an average filter and five different responses from Gaussian and Laplacian-of-Gaussian filters. The bag-of-features is obtained by first computing dense SIFT descriptors [18] in the superpixel region at three different scales and then encoding the descriptors into a dictionary of 400 visual words learned by K -means clustering. We concatenate the appearance-based descriptor and the bag-of-features into a 451-D vector, noted as s_i^t for the t -th superpixel of the i -th image.

Once the feature vector is extracted, a superpixel classifier is built by multiclass SVM [19, 20] with a linear kernel trained over the MSRC-21 dataset. Please note that this dataset is a separate dataset from the action dataset and that object detectors will not be re-trained. Once the classifier is trained, for every measurement, s_i^t , we compute and collect the probability scores of all 23 classes as a feature vector, x_i^t :

$$[x_i^t = p(k|s_i^t) \propto \exp(w_k^T s_i^t)], \forall k \quad (1)$$

where w_k notes the k -th class' parameter vector of the SVM multiclass classifier. Such a vector of posterior probabilities will be later used as the superpixel's measurement for action classification, exploiting the semantic of the object classes and reducing the measurement dimensionality from 451-D to 23-D.

3. LATENT STRUCTURAL SVM

For action classification, we wish to learn the following linear prediction function:

$$(\bar{y}, \bar{h}) = f_w(x) = \operatorname{argmax}_{y, h} [w^T \psi(x, h, y)] \quad (2)$$

where $\psi(x, h, y)$ is a generalised feature function computing a combined map over measurements x , states h and class y , and w is a corresponding parameter vector. The action class and superpixel states are predicted jointly and typically only the predicted class, \bar{y} , is retained. For parameter estimation, we adopt the well-established latent structural SVM framework [21]. This is a regularised minimum-risk framework guaranteed to provide a local optimum for structural models with latent variables. Its learning objective:

$$\begin{aligned} w^* = \operatorname{argmin}_{w, \xi_{1:N}} & \frac{1}{2} \|w\|^2 + C \sum_{i=1}^N \xi_i \\ \text{s.t. } & w^T \Psi(x_i, h_i^*, y_i) - w^T \Psi(x_i, h, y) \geq 1 - \xi_i \\ & \forall \{y, h\} \neq \{y_i, h_i^*\} \end{aligned} \quad (3)$$

$$h_i^* = \operatorname{argmax}_h w^{*T} \Psi(x_i, h, y_i) \quad (4)$$

is an iterative objective that alternates between the constrained optimisation in (3), performed using the current values for latent variables h_i^* , and a new assignment for h_i^* (4) from updated model w^* . Implementation requires a number of design choices including the definition of a suitable feature function, $\psi(x, h, y)$, the initialisation of latent variables h , and efficient algorithms for inference and augmented inference. These components are presented in the following sub-sections.

The learning procedure in (3-4) can be initialised by either an arbitrary vector w^* in (4) or an arbitrary assignment for the h_i^* in (3). Given that we have trained a multiclass superpixel classifier to obtain the feature set, the most natural choice is to initialise the states with the prediction from this classifier:

$$h_{i_{init}}^{*t} = \operatorname{argmax}_k [p(k|s_i^t)] \quad (5)$$

The above is equivalent to initialising the states with individual predictions, delegating the discovery of correlations to the training stage.

3.1. Feature function and score function

The features in feature function $\psi(x, h, y)$ reflect the topology of the graphical model that includes an edge between each superpixel's measurement and its state variable, a fully-connected graph amongst states, and an edge between each state and the action class. In detail, $\psi(x, h, y)$ breaks into:

- *measurement features*, $\varphi(x^t, h^t = j)$: these features map the measurement vector of the t -th superpixel, x^t (dimensionality: K) to its state, h^t (possible values: $\{1 \dots K\}$). The size of this feature vector is K^2 and, given $h^t = j$, it consists of x^t starting at index $K(j-1)$ and zero-padding elsewhere;
- *state features*, $\theta(h^t = j, h^u = k)$: these features report the co-occurrence of states $h^t = j$ and $h^u = k$. The size of this feature vector is again K^2 , with a value 1 at index $K(j-1) + k - 1$ and zeros elsewhere;
- *class features*, $\phi(y = b \in \{0, 1\}, h^t = j)$: these features report the co-occurrence of action class $y = b$ and state $h^t = j$. The size of this feature vector is $2K$, with a value 1 at index $bK + j - 1$ and zeros elsewhere.

We refer the reader to [20] for further details on feature maps. Given such feature vectors, the score function in (2) is computed as:

$$\begin{aligned} w^T \psi(x, h, y) = & \sum_{t=1}^T w_\varphi^T \varphi(x^t, h^t) + \sum_{t=1}^T \sum_{\substack{u=1, \\ u \neq t}}^T w_\theta^T \theta(h^t, h^u) \\ & + \sum_{t=1}^T w_\phi^T \phi(y, h^t) \end{aligned} \quad (6)$$

where $w^T = [w_\varphi^T w_\theta^T w_\phi^T]$ is the concatenation of the parameter vectors for the corresponding features.

3.2. Loss-augmented inference

Following the structured learning approach of [21], a fundamental step in the learning procedure is the computation of a loss-augmented version of the inference. As loss function,

we simply use the 0-1 loss function:

$$\Delta(y^{gt}, y) = \begin{cases} 0 & \text{if } y \neq y^{gt} \\ 1 & \text{if } y = y^{gt} \end{cases}$$

where y^{gt} represents the ground-truth label. With this choice, it can easily be seen that the loss-augmented inference:

$$(\bar{y}, \bar{h}) = \operatorname{argmax}_{y, h} [w^T \psi(x, h, y) + \Delta(y^{gt}, y)] \quad (7)$$

is equivalent to the standard inference in (2) with the addition of a unit score over the incorrect class.

4. EXPERIMENTAL RESULTS

We have evaluated the proposed approach on the most challenging static action recognition dataset released to date, Stanford 40 Actions [22]. This dataset contains images of humans performing 40 different classes of actions, including visually-challenging cases such as “fixing a bike” versus “riding a bike” or “phoning” versus “texting message”; the full class list is provided in the annotation of Fig. 2. The number of samples per class varies between 180 and 300, for a total of 9,532 images. A standard training/test split is made available by the authors on their website, selecting 100 images from each class for training and leaving the remaining for testing.

Training single-class classifiers, also referred to as detectors, using all the available training samples leads to a very unbalanced set over the positive and negative classes (100 and 3,900 samples, respectively). In the case of a conventional SVM objective as in (3), this biases the prediction function towards negative predictions. For this reason and in order to save learning time, we decided to sub-sample the negative training samples of each classifier by randomly choosing 5 images from each of its 39 negative classes. Parameter C in (3) was set to 1.

When measuring the performance of a detector, it would be trivial to achieve high overall accuracies by always predicting the negative class. Therefore, a dataset like Stanford 40 Actions requires measuring the accuracy for the positive and negative classes separately, or providing measures such as precision and recall, average precision at various levels of recall, or similar. In this work, we decided to report the accuracy for the positive and negative classes as the main figure, as shown in Fig. 2. The mean accuracy over all the 40 classes proved 74.06% for the positive class and 88.50% for the negative class. These results prove that the individual detectors are well balanced over positive and negative predictions and show that the overall accuracy is much higher, for




Fig. 2. Achieved accuracy over the positive and negative class for all classes in the Stanford 40 Actions dataset.

instance, than that recently reported in [23] (55.93%). For the individual classes, some classes reach very high accuracy over both positive and negative samples. For instance, class “rowing a boat” achieves 93.90% and 96.15% accuracy, respectively. Other classes show significant missed detections: for instance, class “running” achieves only 48.99% accuracy over the positive samples, arguably as it is hard to recognise a running action from a static frame. Overall, these results seem even more remarkable considering that the 23 object detectors were trained on a completely separate dataset (no re-training was performed on the action dataset) and from classes mostly unrelated with the objects portrayed in Stanford 40 Actions (such as bikes, phones, cameras, telescopes, microscopes and several others).

5. CONCLUSION

In this paper, we have proposed an approach to action recognition in still images leveraging an intermediate superpixel representation of the image. The approach consists of two stages: in a first stage, the image is segmented into a set of superpixels and an array of trained object detectors is applied to each superpixel to extract a vector of detector scores. In a second stage, the score vectors are used as measurements in a graphical model that jointly predicts the superpixels’ classes together with the action class. Experiments conducted over the highly challenging Stanford 40 Actions dataset have resulted in a remarkable accuracy of 74.06% for the positive

class and 88.50% for the negative class averaged over the 40 action classifiers. These results give evidence to the existence of a useful relationship between the classes of the image superpixels and that of the main action. Possible ways to further improve the performance of the proposed model would be to adopt a larger, more universal set of object detectors, or a set of detectors tuned in the specific object classes of given action datasets.

6. REFERENCES

- [1] V. Delaitre, I. Laptev, and J. Sivic, "Recognizing human actions in still images: a study of bag-of-features and part-based representations," in *Proceedings of the British Machine Vision Conference*, 2010, pp. 1–11.
- [2] Nazli Ikizler, R. Gokberk Cinbis, Selen Pehlivan, and Pinar Duygulu, "Recognizing actions from still images," in *Proceedings of the 19th International Conference on Pattern Recognition, 2008, ICPR 2008*, 2008, pp. 1–4.
- [3] I. Laptev, "On space-time interest points," *International Journal of Computer Vision*, vol. 64, no. 2, pp. 107–123, 2005.
- [4] Navneet Dalal and Bill Triggs, "Histograms of oriented gradients for human detection," in *International Conference on Computer Vision & Pattern Recognition*, June 2005, vol. 2, pp. 886–893.
- [5] D.G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [6] Aude Oliva and Antonio Torralba, "Modeling the shape of the scene: A holistic representation of the spatial envelope," *Int. J. Comput. Vision*, vol. 42, no. 3, pp. 145–175, May 2001.
- [7] N. Ikizler-Cinbis, R.G. Cinbis, and S. Sclaroff, "Learning actions from the web," in *2009 IEEE 12th International Conference on Computer Vision*. IEEE, 2009, pp. 995–1002.
- [8] A. Gupta, A. Kembhavi, and L.S. Davis, "Observing human-object interactions: Using spatial and functional compatibility for recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 10, pp. 1775–1789, oct. 2009.
- [9] Bangpeng Yao and Li Fei-Fei, "Grouplet: A structured image representation for recognizing human and object interactions," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010*, 2010, pp. 9–16.
- [10] Weilong Yang, Yang Wang, and Greg Mori, "Recognizing human actions from still images with latent poses," in *CVPR*, 2010, pp. 2030–2037.
- [11] Bangpeng Yao, Xiaoye Jiang, Aditya Khosla, Andy Lai Lin, Leonidas J. Guibas, and Li Fei-Fei, "Action recognition by bases of action attributes and parts," in *International Conference on Computer Vision (ICCV)*, 2011, pp. 1331–1338.
- [12] M. Fischler and R. Elschlager, "The representation and matching of pictorial structures," *IEEE Transactions on Computers*, vol. 22, no. 1, pp. 67–92, Jan. 1973.
- [13] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, "Object detection with discriminatively trained part-based models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 9, pp. 1627–1645, Sept. 2010.
- [14] Guodong Guo and Alice Lai, "A survey on still image based human action recognition," *Pattern Recognition*, vol. in press, no. 0, pp. –, 2014.
- [15] Pedro F Felzenszwalb and Daniel P Huttenlocher, "Efficient graph-based image segmentation," *International Journal of Computer Vision*, vol. 59, no. 2, pp. 167–181, 2004.
- [16] Deli Pei, Zhenguo Li, Rongrong Ji, and Fuchun Sun, "Efficient semantic image segmentation with multi-class ranking prior," *Computer Vision and Image Understanding*, 2013.
- [17] A. Criminisi, "Microsoft research cambridge object recognition image database, <http://research.microsoft.com/en-us/projects/objectclassrecognition/>," 2004.
- [18] A. Vedaldi and B. Fulkerson, "VLFeat: An open and portable library of computer vision algorithms," <http://www.vlfeat.org/>, 2008.
- [19] Koby Crammer and Yoram Singer, "On the algorithmic implementation of multiclass kernel-based vector machines," *The Journal of Machine Learning Research*, vol. 2, pp. 265–292, 2002.
- [20] I. Tsochantaris, T. Joachims, T. Hofmann, and Y. Altun, "Large margin methods for structured and interdependent output variables," *JMLR*, vol. 6, pp. 1453–1484, 2005.
- [21] Chun-Nam John Yu and Thorsten Joachims, "Learning structural svms with latent variables," in *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 2009, pp. 1169–1176.
- [22] Bangpeng Yao, Xiaoye Jiang, Aditya Khosla, Andy Lai Lin, Leonidas Guibas, and Li Fei-Fei, "Human action recognition by learning bases of action attributes and parts," in *Computer Vision (ICCV), 2011 IEEE International Conference on*. IEEE, 2011, pp. 1331–1338.

- [23] Fadime Sener, Cagdas Bas, and Nazli Ikizler-Cinbis, “On recognizing actions in still images via multiple features,” in *Computer Vision ECCV 2012. Workshops and Demonstrations*, 2012, vol. 7585 of *Lecture Notes in Computer Science*, pp. 263–272.